Local Matching Dynamics in Social Networks*

Martin Hoefer

Dept. of Computer Science, RWTH Aachen University, Germany
mhoefer@cs.rwth-aachen.de

Abstract. We study stable marriage and roommates problems in graphs with locality constraints. Each player is a node in a social network and has an incentive to match with other players. The value of a match is specified by an edge weight. Players explore possible matches only based on their current neighborhood. We study convergence of natural better-response dynamics that converge to *locally stable matchings* – matchings that allow no incentive to deviate with respect to their imposed information structure in the social network. For every starting state we construct in polynomial time a sequence of polynomially many better-response moves to a locally stable matching. However, for a large class of oblivious dynamics including random and concurrent better-response the convergence time turns out to be exponential. In contrast, convergence time becomes polynomial if we allow the players to have a small amount of random memory, even for many-to-many matchings and more general notions of neighborhood.

1 Introduction

Matching problems are at the basis of many important assignment and allocation tasks encountered in economics and computer science. A prominent model for these scenarios are *stable matching* problems [13], as they capture the aspect of rationality and distributed control that is inherent in many assignment problems today. A variety of allocation problems in markets can successfully be analyzed within the context of two-sided stable matching, e.g., the assignment of jobs to workers [5, 14], organs to patients [19], or general buyers to sellers. In addition, stable marriage problems are an interesting approach to model a variety of distributed resource allocation problems in networks [3, 12, 18].

In this paper, we examine a dynamic variant of stable matching for collaboration in (social) networks without central coordination. Players are rational agents that are looking for partners for a joint activity or relationship, such as, e.g., to do sports, write a research paper, share an office, exchange data etc. Such problems are of fundamental interest in economics and sociology, and they serve as basic coordination tasks in computer networks with distributed control. We can capture these problems within the stable roommates problem, an extension of stable marriage that allows arbitrary pairs of players to be matched. A crucial aspect of ordinary stable marriage and roommates problems is that every player

* Supported by DFG grant Ho 3831/3-1.

L. Aceto, M. Henzinger, and J. Sgall (Eds.): ICALP 2011, Part II, LNCS 6756, pp. 113–124, 2011.

knows the complete player set and can match arbitrarily. In contrast, for many of the examples above, we would not expect a player to be able to, e.g., write a research paper with any other player instantaneously. Instead, there are often restrictions in terms of knowledge and information that allow certain players to match up easily, while others need to get to know each other first before they can engage in a joint project. We incorporate this aspect by assuming that players are nodes in a static social network of existing *social links*. Each player strives to build a *matching edge* to another player. The network defines a dynamic information structure over the players, where we use a standard idea from social network theory called *triadic closure*: If a knows b and b knows c, then a and c are likely to meet and thus can engage in a joint project. When matching a to the 2-hop neighbor c, both players engage in a joint project and thus become more familiar with each other. In this case, triadic closure suggests that a learns about all direct neighbors of c, which can allow him to find a better matching partner among those players. More formally, at any point in time, we assume that each player can match only to *accessible* players, that is, players in the 2-hop neighborhood in the graph composed of social links and currently existing matching edges.

Traditionally, in the stable marriage problem we have sets of men and women, and each man (woman) has a full preference list over all women (men). Each man (woman) can be matched to exactly one woman (man), where all players would rather be matched than unmatched. A *blocking pair* is a pair of a man and a woman such that both prefer this match to their currently assigned partners (if any), and a *stable matching* is a matching that allows no blocking pair. It is well-known that in this case a stable matching always exists and can be found in polynomial time using the classic Gale-Shapley algorithm [11]. This does not hold for the extension to the stable roommates problem, where simple examples without stable matching exist. In this paper, we will focus on the prominent class of *correlated* [1, 2] (also called acyclic [18]) preferences. For the correlated stable roommates problem existence of a stable matching is guaranteed [1].

Contribution For many of the assignment and matching tasks in (social) networks that motivate our study, there is an inherent lack of information and coordination. Hence, we are interested in distributed dynamics that allow players with locality restrictions to reach a stable matching quickly. In particular, we consider convergence to *locally stable matchings* – matchings that allow no blocking pair of accessible players. We will focus on variants of sequential *best-response* and *better-response dynamics*, where in each round a blocking pair of accessible players (or *local blocking pair*) is allowed to deviate to establish their joint edge. It follows directly from results in, e.g., [1, 2] that our games are potential games, and thus all such dynamics are guaranteed to converge to a locally stable matching. This holds even for more general variants, in which each player can build up to $k > 1$ matching edges, or each player has access to all players within $\ell > 2$ hops in the graph.

After a formal definition of our model in Section 2, we show in Section 3 that in our basic game with $k = 1$ matching edges and lookahead $\ell = 2$ per player we can always achieve fast convergence – for every game and every starting state there is a sequence of polynomially many better-response moves to a locally stable matching. While this shows that locally stable matchings are achievable, the sequence heavily relies on global knowledge of the graph. In contrast, oblivious dynamics without such knowledge like random or concurrent better-response need an exponential number of steps. When we turn to more general games with $k > 1$ or $\ell > 2$, there are even games and starting states such that *every* sequence of better-response moves to a locally stable matching is exponentially long. For the special case of stable marriage, we show that for general preference relations there are states, from which convergence might never be achieved. We also show improved results for special cases of the problem, especially for one-sided social networks and the job-market model of [5].

Perhaps surprisingly, instead of the usual structural aspects in social networks such as, e.g., low diameter or power-law degree distribution, a natural aspect resulting in polynomial time convergence is memory. In Section 5 we consider the case that every polynomial number of rounds each player remembers (uniformly at random) one of the players he had been matched to before. This seemingly small but powerful adjustment allows to show polynomial-time convergence for a variety of dynamics, for arbitrary $k \geq 1$ and $\ell \geq 2$. Finally, we also briefly touch upon the case when memory is considered as a cache in which a bounded number of good or recent matches are stored. Here we can again show an exponential lower bound for standard eviction strategies such as FIFO or LRU.

Related Work. There has been an enormous research interest in stable marriage and roommates problems over the last decades, especially in many-to-one matchings and preference lists with ties [6, 9, 14, 21]. For a general introduction to the topic, see standard textbooks [13, 21].

Recent theoretical work on convergence issues in ordinary stable marriage has focused on better-response dynamics, in which players sequentially deviate to blocking pairs. It is known that for stable marriage these dynamics can cycle [17]. On the other hand, there is always a sequence of polynomially many moves to a stable matching [20]. However, if the ordering of moves is chosen uniformly at random at each step, convergence time is exponential [2]. A prominent case with numerous applications [1, 4, 12, 18], in which fast convergence is achieved even by random dynamics, is correlated or weighted stable matching where each matched pair generates a benefit (or edge weight) for the incident players and the preferences of the players are ordered with respect to edge benefit [1, 18].

Local aspects of stable matching are of interest in distributed computing, e.g., communication complexity of distributed algorithms [16]. A localized version of stable marriage is analyzed in [10], where men and women are nodes in a graph and can only match to adjacent women or men, resp. Each player can only exchange messages with their neighbors and the goal is to design a local algorithm that computes an "almost" stable matching. Similar approaches to almost stable matchings in decentralized settings include, e.g., [7]. In addition, there exist online [15] and parallel algorithms [8] for stable marriage.

Our model of locality is similar to Arcaute and Vassilvitskii [5], who consider locally stable matchings in a specialized case of stable marriage. In their job-market game, there are firms that strive to hire workers. Social links exist only among workers, and each firm can match to k workers, but each worker only to one firm. They show that best-response dynamics converge almost surely and show several characterization results for locally stable matchings and the number of isolated firms after a run of a local variant of the Gale-Shapley algorithm. In this paper we greatly extend their results on convergence of dynamics.

2 Model and Initial Results

In our model we are given a social *network* $N = (V, L)$, where V is a set of n *players* and $L \subseteq V \times V$ is a set of undirected and fixed social *links*. In addition, we have a set E of undirected *potential matching edges*, where we denote $m = |E|$. Each edge $e \in E$ has a *benefit* $b(e) > 0$. A *state* $M \subseteq E$ of the game contains for each player at most k incident matching edges, i.e., each player can be involved in up to $k \geq 1$ matching edges simultaneously. Unless specified otherwise we will assume throughout that $k = 1$. The *utility* or *welfare* of a player u in state M is $\sum_{\{u,v\} \in M} b(\{u, v\})$ if he is matched to at least one player and 0 otherwise. The restriction of E to a subset of all edges will be mostly for convenience and clarity of presentation. Our lower bounds can be adjusted to allow every pair as matching edge using minor technical adjustments that blow up the network using dummy players and for $e \notin E$ use benefits that are either extremely tiny (to keep an edge from becoming a blocking pair) or extremely large (to "hard-wire" an edge and thereby change the matching incentives). Details are left for the full version of this paper.

In a state M two players u and v are *accessible* if there is a path of length at most ℓ in the *access graph* $G = (V, L \cup M)$. We call ℓ the *lookahead* of the game and, unless stated otherwise, focus on the case of triadic closure and $\ell = 2$. An edge $e = \{u, v\} \in E$ is called a *local blocking pair* for M if u and v are accessible, and if for each of u and v either (a) the player has less than k matching edges in M or (b) at least one incident edge $e' \in M$ has $b(e') < b(e)$. Hence, for a local blocking pair the accessible players both strictly increase their welfare by either adding e or replacing e' by e. In the latter case, the replaced edges are removed from M. We call $b(e)$ the benefit of the local blocking pair. A *locally stable matching* is a state M for which there is no local blocking pair.

We consider round-based improvement dynamics that lead to locally stable matchings. In a *local improvement move* we pick a local blocking pair and allow the involved players to deviate to their joint edge, thereby potentially removing other edges. We consider sequential processes that in every round implement a local improvement move. For *(random) best-response dynamics*, we pick in each round deterministically (uniformly at random) a local blocking pair of maximum benefit. As *(random) better-response dynamics* we term all sequential dynamics that in each round pick one local blocking pair in a deterministic (uniformly at random) way. Finally, for *concurrent better- (best-)response dynamics* we assume that every player picks uniformly at random from the local blocking pairs he is involved in (and that are of maximum benefit among those available for him). For every local blocking pair that is chosen by both incident players the corresponding edge is built concurrently.

In addition, we also consider better-response dynamics with memory. For a dynamics with *random memory* we assume that each player at some point recalls a player that he had been matched to before. In particular, let M_v be the set of players that $v \in V$ has been matched with at some point during the history of play. We assume that, in expectation, every T rounds a player v remembers a player u chosen uniformly at random from M_v, and u and v become temporarily accessible in this round. For dynamics with *cache memory* we assume that each player has a cache in which he can keep up to r players previously matched to him. A pair of players then is accessible if and only if they are currently at hop distance at most ℓ in G or one player is present in the cache of the other player.

Our games reduce to the ordinary correlated stable roommates problem for $L = V \times V$. Thus, every ordinary stable matching is a locally stable matching for every network N. Note that an ordinary stable matching can be computed in time $O(n \log n)$ by repeated addition of an edge for a blocking pair with maximum benefit. Hence, centralized computation of a locally stable matching is trivially in P, for every $k \geq 1$ and every $\ell \geq 2$. For ordinary correlated stable roommates even random best-response dynamics converge in polynomial time [2].

3 Convergence of Better-Response Dynamics

In this section we consider the duration of sequential and concurrent improvement dynamics. Even when we restrict to accessible players, a new edge built due to a local blocking pair destroys only edges of strictly smaller benefit. This implies the existence of a lexicographical potential function. Hence, both sequential and concurrent better-response dynamics are always guaranteed to converge to a locally stable matching. Moreover, for our standard case with $k = 1$ and $\ell = 2$ the first result shows that we can always achieve fast convergence.

Theorem 1. *For any state there is a sequence of $O(n \cdot m^2)$ local improvement moves that lead to a locally stable matching. The sequence can be computed in polynomial time.*

Proof. Consider a game with a given network N and a state specified by a set M of existing matching edges. A local blocking pair with edge e falls in one of two categories (1) the players are at distance at most 2 in N or (2) the players are connected via one existing matching edge e' and one link from L. If e falls in category (2), edge e' gets destroyed, and we can think of the edge moving from e' to e. This is the motivation for our main tool in this proof, the *edge movement graph* G_{mov}. This vertex set of this graph is E. Each vertex $\{u, v\} \in E$ has a corresponding vertex weight $b(\{u, v\})$. If u and v are at distance at most 2 in N, their vertex in G_{mov} is called *starting point*. We denote the set of all starting points by S.

There are two kinds of edges in G_{mov}, *movement edges* and *domination edges*. For every triple of players $u, v, w \in V$ we introduce a directed movement edge in

G_{mov} from $\{u,v\}$ to $\{u,w\}$ when $\{u,v\},\{u,w\} \in E$, $\{v,w\} \in L$ and $b(\{u,v\}) < b(\{u,w\})$. When edge $\{u,v\}$ exists, then u and w get accessible. If there is no other matching edge in the system, $\{u,w\}$ becomes a local blocking pair which is expressed by the movement edge. Note that movement edges induce a DAG. Domination edges describe the fact that potentially the existence of one matching edge prohibits creation of another one. For all pairs $\{u,v\}$ and $\{u,w\}$ in G_{mov} we introduce a directed domination edge from $\{u,v\}$ to $\{u,w\}$ when $b(\{u,v\}) \geq b(\{u,w\})$. In this case $\{u,w\}$ is *dominated* by $\{u,v\}$. If the $b(\{u,v\}) > b(\{u,w\})$, $\{u,w\}$ is *strictly dominated* by $\{u,v\}$.

Consider the subgraph of G_{mov} that is reachable from $S \cup M$ via movement edges. If a pair is not in this subgraph, there is no sequence of local improvement moves starting from M that establishes an edge between these players. We prune the graph and consider only the subgraph reachable from $S \cup M$.

A state M of the game can be seen as a marking of the vertices in G_{mov} that correspond to edges in M. A local improvement move from M can only happen if some marked vertex p has an outgoing movement edge to another vertex p' (as $k = 1$, p' must be unmarked). This represents a feasible local improvement move only if p' is currently *undominated*, i.e., has no incoming domination edge from a marked vertex. We will describe how to migrate the markings along movement edges to reach a locally stable matching in a polynomial number of rounds.

Phase 1: In phase 1 we move existing markings without introducing new ones. As long as it is possible, we arbitrarily move an existing marking along a movement edge to an undominated vertex one at a time. Note that the set of dominated vertices changes in every round. In particular, a marking is deleted if it becomes dominated in the next round (because a dominating neighbor with higher benefit becomes marked). Thus, in this process the number of markings only decreases. In addition, for each marking, the vertex weight of the location only increases. Due to acyclicity, a particular marking can visit each node of G_{mov} only once, thus the phase ends after at most $O(n \cdot m)$ many rounds.

Phase 2: In phase 2 we try to improve existing markings even further by introducing new ones and moving them via undominated paths to strictly dominating vertices. In particular, for a marked vertex $\{u,v\}$ in G_{mov} we do the following. We drop all currently dominated vertices from consideration. Now we try to find a path of movement edges from a starting point to a vertex that strictly dominates $\{u,v\}$. If there is such a path, we can introduce a new matching edge via a new marking at the starting point and move it along the path to the dominating vertex. Due to the fact that none of the path vertices are dominated, all the moves are local improvement moves in the game. All markings that become dominated during this process are removed. This also includes in the last step our original marking at $\{u,v\}$, which we can think of as moving to the dominating vertex. After this, we try to improve all markings by a restart of phase 1. We keep executing this procedure until this kind of improvement is impossible for every remaining marking.

Phase 2 can be seen as an extension of phase 1. Overall, we keep decreasing the number of markings, and each surviving marking is increased in terms of vertex weight. However, each such increase might now require $O(m)$ steps, which increases the number of rounds to at most $O(n \cdot m^2)$.

Phase 3: After phase 2, none of the remaining markings can be (re)moved, neither by moving the marking along a movement edge to another undominated vertex, nor by a sequence of moves that lead to creation of a marking at a dominating vertex (verify that otherwise phase 2 would not have ended). Hence, these edges are stable, and we call the incident players *stabilized.* They will not become part of any local blocking pair in the remaining process. We drop all these players from the game and adjust G_{mov} by dropping all vertices including at least one stabilized player. Finally, we now iteratively construct the matching edge of largest possible benefit until no more edge can be constructed. In particular, we consider the reduced G_{mov} (which is now completely unmarked) and find a vertex with largest benefit that is reachable from a starting point. We then establish the corresponding edge by moving a new marking along the path. There is no path to any edge with strictly larger benefit, and no player will get an incentive to remove this edge. In particular, after removing edges and incident players, no such path will become possible at a later point in the process. Hence, iterative application of this final step completes the locally stable matching. This phase terminates after $O(n \cdot m)$ rounds in total.

The construction and adjustment of G_{mov} can be trivially be done in polynomial time. For finding the sequences of local improvement moves we essentially require only algorithms for DFS or BFS in DAGs. □

The computation of the short sequence in the previous theorem relies heavily on identification of proper undominated paths in the edge movement graph. If we instead consider random or concurrent dynamics that do not rely on such global information, convergence time can become exponential.

Theorem 2. *For every $b \in \mathbb{N}$ there is a game with $n \in \Theta(b)$, in which (random) best-response and random and concurrent better-response dynamics starting from the state $M = \emptyset$ need $\Omega(1.5^b)$ steps in expectation to converge to a locally stable matching.*

Proof. We construct our game based on a structure we call an "edge trap". We concatenate b edge traps as shown in Fig. 1. In a trap i there is a starting edge e_i. We assume that initially there is no matching edge and lower bound the number of times that e_i must be created to lead to construction of e_{i+1}. Upon creation of e_i, we assume there are exactly two local blocking pairs that both destroy e_i – either create e^u or e^l, which implies $b(e_i) < \min\{b(e^u), b(e^l)\}$. If e^u is formed, $\{w_i, v_{i+1}\}$ becomes a local blocking pair with e^c destroying edge e^u (i.e., $b(e^u) < b(e^c)$). At this point, both w_i and v_{i+1} are happy with their matching choice. Now suppose e_i is created again, then player w_i will not form e^u. In this case, v_i matches to u_{i+1} via e^l. Now v_{i+1} has an incentive to drop e^c and match with u_{i+1}, because $b(e^c) < b(e_{i+1})$.

By sequential concatenation of edge traps we can make the dynamics simulate a counter with b bits. In particular, we assume there are b traps attached sequentially, where in the first trap players u_1 and v_1 are connected with a path

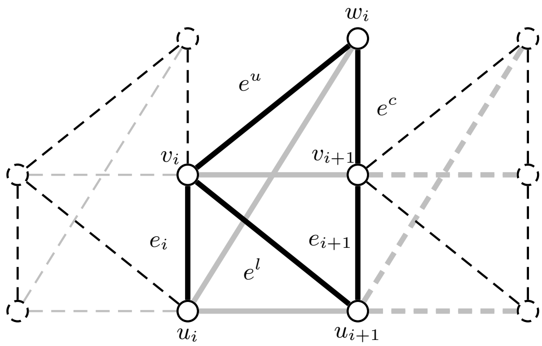

Fig. 1. Structure of edge traps in the lower bound of Theorem 2. Social links are indicated in gray, possible matching edges are drawn in black.

of distance 2 using a dummy player. For the last pair of nodes u_{b+1} and v_{b+1} we assume there is also the final edge e_{b+1}. Bit i is set if and only if edge e^c in trap i is created. We start the dynamics with the empty matching $M = \emptyset$, so the counter is 0. Note that N is a tree with a long path and satellite nodes.

First consider best-response dynamics, for which we set $b(e^u) > b(e^l)$ in every trap. Then creation of e_1 implies creation of e^c in the first trap, i.e., increase of the counter by 1. At this point, the edge is trapped, and the only local improving move is to create e_1 again. This leads to destruction of e^c in the first trap, creation of e^c in the second trap, and thus an increase in the bit counter of 1. Continuing in this fashion we see that every creation of e_1 leads to a state that represents an increase of the bit counter by 1. Thus, to create e_{b+1}, the edge of largest benefit, the dynamics needs $\Omega(2^b)$ many creations of e_1. Note that in each step there is a unique local blocking pair of maximum benefit. This proves the lower bound for both deterministic and random best-response dynamics.

Note that for each state reached during the dynamics, only one local blocking pair except (u_1, v_1) can apply their deviation. In particular, e^u and e^l cannot be created simultaneously. Hence, as long as the creation of e^u is sufficiently likely in every trap, we can trap enough edges that are subsequently destroyed and show a similar lower bound. This is the case, in particular, for random and concurrent better-response dynamics. Note that every locally stable matching must contain the edge e_{b+1}. We now bound the number of times we have to create e_1 until e_{b+1} forms. Consider the first trap, and suppose edge e_1 is created. We now follow this edge moving through the trap. With probability at most 1/2, the edge moves directly to e^l and arrives at e_2. With probability at least 1/2, the edge gets stuck in e^c, which implies that the edge is destroyed and the next edge created at e_1 arrives at e_2 with probability 1. Thus, to create a single edge at e_2 we have to create e_1 an expected number of 1.5 times. The same is true for e_i and e_{i+1} in every trap i. Thus, due to the sequential nature of the gadget, we need an expected number of $\Omega(1.5^b)$ creations of e_1 to create edge e_{b+1}. □

The previous theorem applies similarly to a wide variety of better-response dynamics. The main property is that whenever e_i exists and both e^u and e^l are available for creation, then the creation of e^u is always at least as likely as that of e^l. Observe that the construction allows to set $b(e^u) < b(e^l)$ or vice versa. This allows to satisfy the property for many dynamics that make "oblivious" choices by picking local blocking pairs based only on their benefits but not their structural position or the history of the dynamics. An even stronger lower bound applies when we increase lookahead or matching edges per player.

Theorem 3. *Let $k \geq 2$ or $\ell \geq 3$, then for every $b \in \mathbb{N}$ there is a game with $n \in \Theta(b \cdot k \cdot \ell)$ and a starting state M such that every sequence of local improvement moves from M to a locally stable matching has length $\Omega(2^b)$.*

By embedding the lower bound structures from the previous proofs into larger graph structures of dummy vertices, we can impose a variety of additional properties on the network N that are often considered in the social network literature. For example, to obtain a small diameter simply add a separate source and connect each vertex from the gadgets via ℓ dummy vertices to the source. As these new vertices have no matching edges, the lower bounds hold accordingly for graphs of diameter $\Theta(\ell)$. Note that for a diameter of exactly ℓ we obtain the ordinary correlated stable roommates problem, for which polynomial time convergence is guaranteed. As mentioned earlier there also exist simple adjustments using dummy vertices and tiny and extremely large benefits to adjust the results to $E = V \times V$ with all possible matching edges.

Corollary 1. *Theorems 2 and 3 continue to hold even if $diam(N) \in \Theta(\ell)$.*

4 Two-Sided Matching and Stable Marriage

In this section we consider the bipartite case of stable marriage. In accordance with [5] we use the interpretation of a set W of "workers" matching to a set F of "firms". We use $n_w = |W|$ and $n_f = |F|$. In general, the social network $N = (F \cup W, L)$ can be arbitrary, and the set of possible matching edges is $E = W \times F$. Each worker (firm) has an arbitrary preference relation over firms (workers). In contrast to ordinary stable marriage, in our localized variant convergence of any better-response dynamics can be impossible.

Theorem 4. *For stable marriage with general preferences, there are games and starting states M such that no locally stable matching can be reached by any sequence of local improvement moves from M, even if N is connected.*

For weighted matching with correlated preferences, Theorem 1 applies and shows existence of a short sequence. A central assumption of [5] is that every worker $w \in W$ has the same preference list over F, and every firm $f \in F$ has the same preference list over W. We will refer to this case as *totally uniform preferences.* As a generalization we consider *worker-uniform* preferences, where we assume that only the preferences of all workers are the same, while firms have arbitrary preferences. *Firm-uniform* preferences are defined accordingly. For totally uniform preferences we can number firms and workers increasingly from least to most preferred in their respective global preference list. For edge $e_{ij} = (w_i, f_j)$ we define a benefit $b(e_{ij}) = j \cdot n_w + i$. Intuitively, here best-response dynamics give preference to local blocking pairs of the most preferred firm, which can be changed to worker by using $b(e_{ij}) = i \cdot n_f + j$ throughout. For worker-uniform preferences we let the numbering of workers be arbitrary. For $e_{ij} = (w_i, f_j)$ we define benefit $b(e_{ij}) = j \cdot n_w + i_j$, when worker w_i is ranked at the i_j-th last position in the preference list of firm f_j. For firm-uniform preferences the same idea can be used by exchanging the roles of firms and workers. This shows that all these cases are classes of correlated stable matching problems. Our first result is that even for totally uniform preferences, convergence of best-response dynamics can be slow.

Theorem 5. *For every $b \in \mathbb{N}$, there is a game with totally uniform preferences, $n_f, n_w \in \Theta(b)$ and a starting state M such that (random) best-response dynamics from M to a locally stable matching take (expected) time $\Omega(2^b)$.*

A case in which such preferences can still lead to quick convergence is in the *job-market game* considered in [5]. Note that for the following theorem we only need worker-uniform preferences.

Theorem 6. *For every job-market game with worker-uniform preferences, in which each firm can create up to $k \geq 1$ matching edges, (random) best-response dynamics converge to a locally stable matching from every starting state in (expected) time $O(n_f \cdot n_w \cdot k)$.*

The result directly extends to random and concurrent better-response dynamics when spend an additional factor of $n_f \cdot n_w$. In contrast, if we consider firm-uniform preferences, a lower bound for best-response dynamics can be shown.

Theorem 7. *For every $b \in \mathbb{N}$ and $k \geq 2$, there is a job-market game with firm-uniform preferences, $n_f \in \Theta(b)$, $n_w \in \Theta(b \cdot k)$ and a starting state M such that (random) best-response dynamics from M to a locally stable matching take (expected) time $\Omega(2^b)$.*

5 Dynamics with Memory

In this section we consider sequential and concurrent better-response dynamics with memory. Our first result is a polynomial time bound for random memory that holds accordingly for random and concurrent better-response dynamics.

Theorem 8. *For every $k, \ell \in \mathbb{N}$ and every game, in which each player can create up to k matching edges and has lookahead ℓ, (random) best-response dynamics with random memory converge to a locally stable matching from every starting state in (expected) time $O(n^2 \cdot m^2 \cdot k \cdot T)$.*

Proof. The main insight is that all the above mentioned dynamics can rely on the information in the random memory to steer the convergence towards a locally stable matching. Let us consider the dynamics in phases. Phase t begins after the dynamics has created t mutually different matching edges at least once (including the ones in the starting state). Let E_t be the set of edges which have been created at least once when entering phase t. During phase t no new edge is created for the first time. The main insight is that in a phase all the dynamics converge in polynomial time. In particular, consider an edge $e \in E_t$ that represents a (global) blocking pair and has maximum benefit. A round in which such an edge is available for creation appears in expectation at most every $n \cdot T$ rounds. When it is available for creation all the dynamics mentioned above will create it with probability at least $\Omega(1/m)$. The deterministic best-response dynamics might create a different edge of maximum benefit, but it has to pick e after at most m such rounds unless e stops being a blocking pair. Thus, after an expected number of $O(n \cdot m \cdot T)$ rounds, e is either created or stops being a blocking pair. In the former case, it will not get removed in phase t again, in the latter case other incident edges from E_t of maximum benefit were created that are not removed in phase t again. Hence, at least one of the incident players has an edge that he is not willing to remove in phase t. By repeated application of this argument, we see that after at most $n \cdot k$ of these steps all blocking pairs have been removed. Thus, after an expected number of $O(n^2 \cdot m \cdot k \cdot T)$ rounds the phase ends in a matching that is stable with respect to E_t. Finally, we note that there are at most m many phases to be considered. □

With random memory no previous matching edge can be completely forgotten during the dynamics. Can we allow players to forget some previous matches and still obtain fast convergence with local dynamics? Towards this end, we here consider two natural examples and show that delayed forgetting of previous matches can be harmful to convergence. In particular, we consider best-response dynamics with cache memory and largest-benefit, FIFO or LRU eviction strategies.

Corollary 2. *If each player keeps only the r best matches in his cache, then for every $b \in \mathbb{N}$ there is a game with $n \in \Theta(b \cdot r)$, in which best-response dynamics starting from the state $M = \emptyset$ need $\Omega(2^b)$ steps to converge to a locally stable matching.*

Theorem 9. *If each player keeps only the r most recent matches in his cache, then for every $b \in \mathbb{N}$ there is a game with $n \in \Theta(b^2 \cdot r)$, in which best-response dynamics starting from the state $M = \emptyset$ need $\Omega(2^b)$ steps to converge to a locally stable matching.*

Acknowledgments

The author would like to thank Siddharth Suri and Heiko Röglin for numerous insightful discussions on the topic.

References

1. Abraham, D., Levavi, A., Manlove, D., O'Malley, G.: The stable roommates problem with globally ranked pairs. Internet Math. 5(4), 493–515 (2008)
2. Ackermann, H., Goldberg, P., Mirrokni, V., Röglin, H., Vöcking, B.: Uncoordinated two-sided matching markets. SIAM J. Comput. 40(1), 92–106 (2011)
3. Akkaya, K., Guneydas, I., Bicak, A.: Autonomous actor positioning in wireless sensor and actor networks using stable-matching. Intl. J. Parallel, Emergent and Distrib. Syst. 25(6), 439–464 (2010)
4. Anshelevich, E., Hoefer, M.: Contribution games in social networks. In: Proc. 18th European Symposium on Algorithms (ESA), vol. 1, pp. 158–169 (2010)
5. Arcaute, E., Vassilvitskii, S.: Social networks and stable matchings in the job market. In: Leonardi, S. (ed.) WINE 2009. LNCS, vol. 5929, pp. 220–231. Springer, Heidelberg (2009)
6. Echenique, F., Oviedo, J.: A theory of stability in many-to-many matching markets. Theoretical Economics 1(2), 233–273 (2006)
7. Eriksson, K., Häggström, O.: Instability of matchings in decentralized markets with various preference structures. Intl. J. Game Theory 36(3-4), 409–420 (2008)
8. Feder, T., Megiddo, N., Plotkin, S.: A sublinear parallel algorithm for stable matching. Theoret. Comput. Sci. 233(1-2), 297–308 (2000)
9. Fleiner, T.: A fixed-point approach to stable matchings and some applications. Math. Oper. Res. 28(1), 103–126 (2003)
10. Floréen, P., Kaski, P., Polishchuk, V., Suomela, J.: Almost stable matchings by truncating the gale-shapley algorithm. Algorithmica 58(1), 102–118 (2010)
11. Gale, D., Shapley, L.: College admissions and the stability of marriage. Amer. Math. Monthly 69(1), 9–15 (1962)
12. Goemans, M., Li, L., Mirrokni, V., Thottan, M.: Market sharing games applied to content distribution in ad-hoc networks. IEEE J. Sel. Area Comm. 24(5), 1020–1033 (2006)
13. Gusfield, D., Irving, R.: The Stable Marriage Problem: Structure and Algorithms. MIT Press, Cambridge (1989)
14. Kelso, A., Crawford, V.: Job matchings, coalition formation and gross substitute. Econometrica 50, 1483–1504 (1982)
15. Khuller, S., Mitchell, S., Vazirani, V.: On-line algorithms for weighted bipartite matching and stable marriages. Theoret. Comput. Sci. 127(2), 255–267 (1994)
16. Kipnis, A., Patt-Shamir, B.: On the complexity of distributed stable matching with small messages. Distributed Computing 23(3), 151–161 (2010)
17. Knuth, D.: Marriages Stables et leurs relations avec d'autres problemes combinatoires. Les Presses de l'Université de Montréal (1976)
18. Mathieu, F.: Self-stabilization in preference-based systems. Peer-to-Peer Netw. Appl. 1(2), 104–121 (2008)
19. Roth, A., Sönmezc, T., Ünverd, M.U.: Pairwise kidney exchange. J. Econ. Theory 125(2), 151–188 (2005)
20. Roth, A., Sotomayor, M.O.: Two-sided Matching: A study in game-theoretic modeling and analysis. Cambridge University Press, Cambridge (1990)
21. Roth, A., Vate, J.H.V.: Random paths to stability in two-sided matching. Econometrica 58(6), 1475–1480 (1990)